arXiv:1611.09119v1 [cs.CV] 28 Nov 2016

Unsupervised Feature Learning With Symmetrically Connected Convolutional Denoising Auto-encoders

Jianfeng Dong[†], **Xiao-Jiao Mao**[†], **Chunhua Shen**[⋆], **Yu-Bin Yang**[†]
[†]State Key Laboratory for Novel Software Technology, Nanjing University, China
[⋆]School of Computer Science, University of Adelaide, Australia

Abstract

Convolutional neural networks (CNNs) have shown their power on many computer vision tasks. However, there are still some limitations, including their dependency to large scale labeled training data and sensitivity to weight initialization. In this paper, we try to address these two problems by proposing a simple yet powerful CNN based denoising auto-encoder which can be trained end-to-end in an unsupervised manner. The network architecture we employ is a fully convolutional auto-encoder with symmetric encoder-decoder connections. The proposed method can not only reconstruct clean images from corrupted ones, but also learn image representation through the reconstruction training. It can further be adapted to find data driven network initialization without using extra training data. Experimental results show that our network can learn good feature from unlabeled data, which can be easily transferred to different high level vision tasks such as image classification and semantic segmentation. The data driven initialization method based on the convolutional auto-encoder is also competitive.

1. Introduction

In recent years, convolutional neural networks (CNNs) have set new state-of-the-arts in many vision tasks, such as image classification, object detection and semantic segmentation. Yet most of them depend on the feature learned from large scale labeled dataset such as ImageNet [23]. While unlabeled data is much easier to get, how to learn good and transferable feature from unlabeled data has become a natural problem.

Denoising auto-encoder [31] is one of the most popular neural network framework for unsupervised learning, which can help to find good network initialization or learn discriminative feature by reconstructing the original data from its corrupted one. The typical architecture is an encoder-decoder network, in which the encoder learns the representation and decoder reconstructs the input from the learned representation. However, the denoising strategy is not quite suitable for training CNNs, even though it has shown great potential in training fully connected networks. The reasons may be: (1) the encoder-decoder architecture normally doubles the depth of the network, thus makes it harder to train end-to-end, (2) convolution usually leads to the loss of image detail, and it is difficult for the small kernels to "remember" rich image detail for reconstrcting the clean image, and (3) for denoising, CNNs learn the noise pattern, which may not be the desired "feature pattern" of the images.

In this paper, we propose a fully convolutional auto-encoder, which has symmetric shortcut connections from encoder to decoder, for unsupervised learning. The fully convolutional architecture can make sure both downsampling and upsampling are learnable. The symmetric skip-layer connections are introduced to address the first two problems. They can ease the training procedure of the double-depth convolutional auto-encoder by directly propagating gradient from top layers to bottom ones. Moreover, they help pass image detail forwardly, which force the middle layers focusing on learning high level abstraction, instead of remembering low level information. To address the third problem, we train the network with carefully selected types of corruption, which shows that the corruption itself is very essential for convolutional auto-encoder.

How to find a suitable network initialization is another challenge for training deep neural networks. Most existing initialization techniques focus on investigating the statistics of a network and making the training process free from gradient vanishing/exploding. However, few of them care about the discriminant of training data. For example, the initialization of a recognition network trained for hand-writing digits should be very different from the one trained for medical images. We show that our network can also learn data-driven network initialization from merely the training data, even when there is no external unlabeled data available.

The remaining content is organized as follows. We provide a brief review of related work in Section 2. We present

the network architecture as well as the unsupervised learning and data driven initialization method in Section 3. Experimental results are provided as analyzed in Section 4.

2. Related work

Unsupervised and weakly supervised feature learning Learning representation from unlabeled data using deep neural network has been long studied. Stacked auto-encoders [2] learn compressed representation by training an encoder network and judge its quality by the reconstruction loss of a decoder network. The training is conducted layer-by-layer. Assuming that good representation is robust to local corruption, denoising auto-encoders [31] are proposed to learn feature by reconstructing clean data from corrupted one. Our unsupervised learning strategy is inspired by denoising auto-encoders, but the architecture is based on deep CNNs and can be trained end-to-end instead of layer-by-layer. Some other CNN-based unsupervised or self-supervised learning methods include [6, 5, 32, 21, 12]. Dosovitskiy *et al*. [6] learned feature by discriminating random augmented image patches. Doersch *et al*. [5] used relative positions of patches from images as supervision. Other supervisions like video tracking[32] and cluster membership [12] are also used to learn feature. A more relevant work is Pathak *et al*. [21]. They learn feature by image inpainting with a CNN based auto-encoder called Context Encoder. In order to learn good features, their network was carefully designed for good image inpainting performance. While our method is more flexible since we can conduct unsupervised feature learning by general image restoration. Moreover, we can incorporate deeper networks and learn better features with the help of shortcut connections.

Weight initialization Initializing a network with small random numbers is usually adequate when the network is shallow. But when the network goes deeper, it may cause problems such as gradient vanishing or exploding, which can be avoided by carefully designing the network architecture or finding a better initialization. Glorot *et al*. [7] proposed a initialization method by properly scaling the random number according to the network statistics, and it can ease the training of network with sigmoid-like activations. He *et al*. [9] proposed a similar strategy for ReLU-activation network. Both of these two initialization methods depend on the network architecture. A data-dependent initialization method was proposed by Krähenbühl *et al*. [15]. They used both data statistics from a set of training data, and network architecture to determine the initialization , which can make all units in the network trained at roughly the same rate. Different from theirs, we use the deep feature learned from the whole raw training data for network initialization.

Shortcut connections As the networks are going deeper, shortcut connection is introduced to tackle the problem of gradient vanishing/exploding during training [28, 10, 11]. In [28], every "inception" block contains a shortcut connection to the next block. In [10, 11], shortcut connections are linked from the input layer all the way up to the output layer, which makes the learned function a "residual" function. For auto-encoder network, symmetric shortcut connections between encoder and decoder are often used to preserve low level information. Ranzato *et al*. [22] pass local information from encoder to decoder of a CNN-based auto-encoder, and recently this architecture is also in semantic segmentation [1]. Our usage of shortcut connections is mainly inspired by Mao *et al*. [18] , as well as Ladder network [30]. Mao *et al*. [18] used convolutional auto-encoder with encoder-decoder connections to deal with low level image restoration problems. Ladder network [30] is a special kind of denoising auto-encoder with noise added to all the feature maps, and it also uses shortcut connections.

3. Symmetrically Connected Convolutional Denoising Auto-encoders

3.1. Architecture

The basic architecture of our network is a fully convolutional auto-encoder. The encoder part is a chain of convolutional layers, and the decoder part is a chain of deconvolutional layers symmetric to the convolutional ones. The corresponding encoder and decoder layers are connected by shortcut connections. Each convolutional/deconvolutional layer is followed by a Batch Normalization [13] layer and a ReLU non-linearity layer. An illustration is shown in Figure 1.

Encoder The encoder acts as feature extractor. Inspired by VGGNet [24], our encoder network mainly contains 3×3 convolutional layers. The down-sampling is conducted by convolution with stride 2 as in [10]. We did not use pooling for down-sampling because as argued in [18], pooling is harmful to the image restoration task. Moreover, fully convolutional network without pooling is already successfully used by several works on image classification [26, 10, 11].

Two specific encoder architectures used in this paper are shown in Table 1. Network-1 is a simple stack of convolution layers and Network-2 is based on VGG-16 [24]. Note that although our encoder is quite simple, one could always incorporate more complex architectures such as [16, 24, 28]. The only thing to do is constructing a suitable decoder. We choose a concise one because we mainly focus on the investigation of unsupervised training instead of the influence of different network architecture.

Decoder The decoder takes features learned by the encoder and reconstructs the "clean" images. We use deconvolution as our decoder unit, which is often referred as learnable up-sampling methods in tasks such as semantic segmentation [20] and image generation [8]. Since we use convolution with stride as learnable down-sampling in en-

	Network-1	Network-2
block 1	conv 3x3,n ... $\times(m_1-1)$	conv 3x3,64
downsample 1	conv 3x3/2, n	conv 3x3/2,64
block 2	conv 3x3,2n ... $\times(m_2-1)$	conv 3x3,128
downsample 2	conv 3x3/2,2n	conv 3x3/2,128
block 3	conv 3x3,4n ... $\times(m_3-1)$	conv 3x3, 256 conv 3x3, 256
downsample 3	conv 3x3/2,4n	conv 3x3/2, 256
block 4	conv 3x3,8n ... $\times(m_4-1)$	conv 3x3, 512 conv 3x3, 512
downsample 4	conv 3x3/2,8n	conv 3x3/2, 512
block 5		conv 3x3, 512 conv 3x3, 512
downsample 5		conv 3x3/2, 512
block 6	global_ave_pool fc,num_class	fc,4096 fc,4096 fc,num_class

Table 1: Network architectures. Network-1 is used for CIFAR-10, CIFAR-100 and STL-10. Network-2 is based on VGG-16 and used for PASCAL VOC.

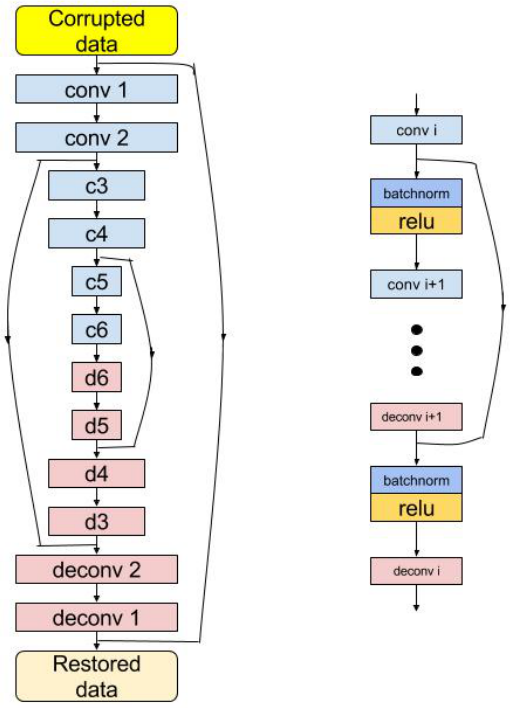

Figure 1: Left: A simple illustration of our network. Right: Shortcut connection

coder, it is important to also make the up-sampling learnable. The layers in the decoder are organized symmetrically to the ones in the encoder. For a specific deconvolutinal layer, the size of its input/output is equal to the output/input of its corresponding convolutional layer. The reason is to achieve pixel-wise correspondence with the use of shortcut connections described later.

Shortcut Connections In our network, symentric encoder-decoder shortcut connections are used to pass signals forwardly. The signals from the shortcut and deconvolutional layers are then added element-wise. We use the shortcut configuration shown in Figure 1. The reasons for using the shortcuts are: Firstly, Mao *et al.* [18] showed that the convolutional encoder-decoder network with shortcut connections can make a great difference in low level image restoration tasks. Denoising auto-encoders assumed that a good representation should be robust to input corruption, and can be used to recover the "clean" input [31]. Therefore, we aim to improve the representation performance of the features by reducing the reconstruction loss with the use of shortcut connections. Secondly, the shortcut connections provide rich image details and help to reconstruct the "clean" input, thus enable the top layers in the encoder to focus on learning better image abstraction instead of trying hard to keep low level image details, which are not quite useful for high level tasks such as classification. At last, building auto-encoders almost doubles the depth of the network. Very deep convolutional auto-encoders based on deep CNNs are hard to train and may easily suffer from problems such as gradient vanishing/exploding. The shortcut connections ease the training by directly back propagating gradients from top layers to bottom layers.

3.2. Unsupervised Learning and Data Driven Initialization

The unsupervised feature learning and evaluation follow a pre-training and fine-tuning pipeline. We first train a denoising convolutional auto-encoder on images without using any labels to learn the features. For a clean image x, the input of the auto-encoder network is its corrupted version $\tilde{x}$, and output is the reconstruction represented as $f(\tilde{x})$. We train the network end-to-end by minimizing the Euclidean loss between x and $f(\tilde{x})$. After the training of denoising network, we make use of the learned features by fine-tuning some of the pre-trained layers together with extra layers for specific tasks such as classification. This pre-training and fine-tuning pipeline has been widely used in transferring neural network based feature from one task to another. A popular pre-training method is to train the network on large

scale labeled datasets. However, our pre-training is purely unsupervised.

Suppose that we pre-train the auto-encoder network on unlabeled dataset A and fine-tune it on dataset B. When $A \subset B$, our pre-training process can be regarded as a data driven weight initialization technique. One may wonder that as all pre-training can be seen as an initialization method, why do we discuss about this case? Firstly, external unlabeled data, which shares the same distribution with the labeled training data, is usually hard to access or even do not exist. However, the training data itself is always at hand. Secondly, since the pre-training are conducted on the data without using the labels, no extra data are used essentially. If the initialization by our pre-training strategy works better than traditional ones (such as random Gaussian) with respect to the overall classification performance, it indicates that our method is able to capture some underlying information from the raw data, which may not be learned in the normal supervised training with simple initialization. Compared to most existing initialization methods that determine the initialization by a specific distribution (such as random Gaussion) or network structure (such as Xavier [7]), our initialization is directly learned from the raw data.

3.3. Corruption type

Traditional denoising auto-encoders only use pixel level corruptions such as pixel-wise Gaussian noise or 0-mask [31]. It is sufficient for fully connected networks, which treat every input (or pixel for image data) equally and do not consider local invariance. However, CNNs tend to model simple local patterns because the convolution operator is locally invariant for its weight sharing and local connectivity properties. In this paper, besides simple pixel level Gaussian noise, we will investigate other types of corruptions. The basic principle to choose the corruptions is: the reconstruction of the corrupted data should be more dependent on large image patches (the whole image if possible) instead of only nearby pixels. In other words, we would like the use of corruption forcing the network to look at a larger part in an image. Three types of corruptions are listed as follows and shown in Figure 2 as well:

Pixel Level Gaussian noise. It is a typical pixel level corruption which shows up in many image denoising and denoising auto-encoder methods. Given an image x, we add random Gaussian noise with 0 mean and standard deviation σ to each pixel uniformly. σ is also called the noise level.

Patch-like Gaussian noise. This type of corruption also adds Gaussian noise to pixels, while a same noise is shared by some adjacent pixels instead of totally uniform. We divide an image into rectangular regions and add a same Gaussian noise to pixels in the same region. Learning to eliminate such corruption is easier than pixel level Gaussian noise since the total amount of noise is far less. The network can also take hints from adjacent pixels, which makes the network trying to look at a larger region.

Block mask. Another kind of corruption is to set some pixels of a image to 0. It is also used in [31] as masking noise. We use a spacial case of the masking noise, in which the pixels dropped out are adjacent. Images with this kind of corruption are used by Pathak *et al*. [21] for their Context Encoder. More specifically, the missing block corruption we use is: taking n square regions randomly from an image, and set pixels in these regions to 0, assuming the image is 0-centred. The length of a side in each square is taken uniformly in $[a, b]$. When n, a and b are not too small, images with this corruption are usually harder to be reconstructed than the ones with pixel level Gaussian noise since it needs more global semantic understanding of the images.

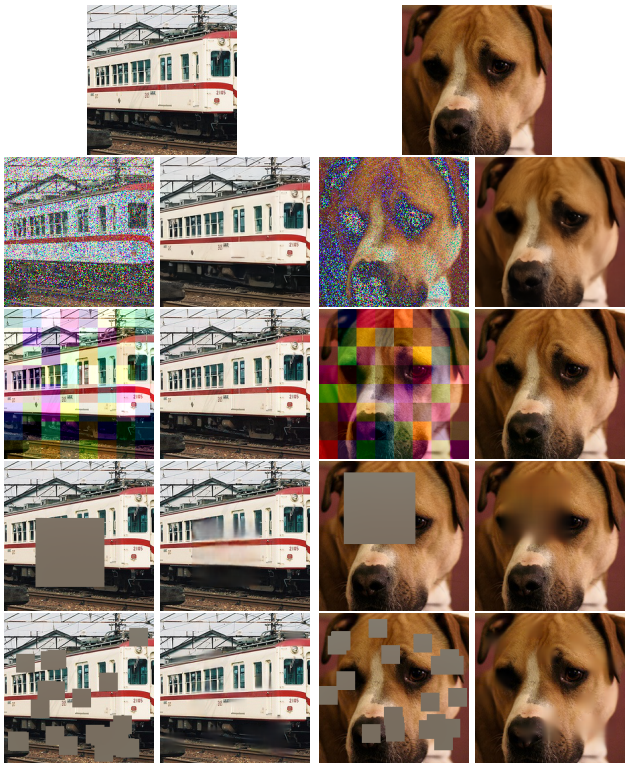

Figure 2: Reconstruction performance on held-out images. The top row indicates clean images. The following rows are 4 types of corrupted images and the reconstructed clean ones. From top to bottom: pixel Gaussian, patch Gaussian, big missing block and small missing blocks.

3.4. Implementation Details

Our method is implemented with MXNet [3]. ADAM [14] is used for optimization. The learning rate is set as 1e-4 at first, and divided by 10 when the loss does not drop. If not specified, the weights are initialized by random Gaussian numbers with 0 mean and standard deviation 0.01.

The detailed encoder architectures are shown in Table 1, and the decoder part is symmetric to the encoder. Shortcut connections are linked every 2 convolutional layers to their corresponding deconvolutional layers, as well as one connection from the input to the output.

We use simple data augmentation for all datasets. For CIFAR-10 and CIFAR-100, we take 29×29 crop are randomly flip it horizontally for each image during training. The crop size is 89×89 for STL-10 and 225×225 for PASCAL VOC. At testing time, central crop is taken for CIFAR-10, CIFAR-100 and STL-10. While for PASCAL VOC classification, we follow [21] to average the results of 10 random crops per test image. All image pixels are 0-centered and normalized. The corruption is added to the image crops at real time before being fed into the network, which means theoretically we have unlimited number of training images for auto-encoder network if the corruption type is not deterministic.

4. Experiments

We firstly evaluate the performance of our initialization strategy. Then, our unsupervised feature learning method is compared with other state-of-the-art methods. And we show the learned feature can be successfully transferred to supervised vision tasks. Afterwards, we investigate multiple types of corruption and shortcut connections. Analytic experiments and some interesting visualization results are presented at last.

4.1. Data-Driven Network Initialization

We evaluate our data driven initialization strategy on 3 different datasets: CIFAR-10, CIFAR-100 and PASCAL VOC 2007 classification. A reconstruction network from corrupted images to clean images is trained on all training images of each dataset without using any label. Then, the weights of the encoder are used to initialize a corresponding classification network. The architecture used for small images is Network-1 in Table 1 with the parameters $m_1 = 5, m_2 = 5, m_3 = 5, m_4 = 0$ (3 times down-sampling instead of 4) and $n = 64$, and VGG based Network-2 for PASCAL VOC. Data augmentation described in Section 3.4 is used during training. We find that pixel Gaussian noise is good enough for pre-training on small images. For larger imgaes in PASCAL VOC, we observe that small block corruption works better.

Table 2 shows the classification accuracy of different network initialization methods and Figure 3 shows the training loss during the first 50 epochs on CIFAR-10. The main observation is that our initialization method outperforms others in both overall accuracy and convergence. Table 3 shows the overall classification accuracy of our network and other state-of-the-art results on CIFAR-10 and CIFAR-100. Our method achieves very competitive results by making use of the proposed data-driven initialization strategy, even if the network itself is very concise comparing to those of other methods.

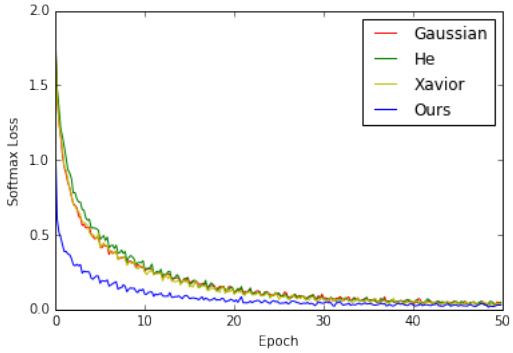

Figure 3: Training loss at first 50 epochs using different initialization methods on CIFAR-10

	CIFAR-10	CIFAR-100	VOC 2007
Gaussian	93.95%	73.89%	67.11%
He *et al*. [9]	93.67%	72.15%	66.71%
Xavior[7]	93.60%	72.9%	67.27%
Ours	**94.93%**	**75.41%**	**69.87%**

Table 2: Comparisons with other initialization methods: CIFAR-10, CIFAR-100 and PASCAL VOC 2007 classification results.

4.2. Unsupervised Feature Learning

In this section we evaluate our unsupervised feature learning method on image classification and semantic segmentation.

4.2.1 CIFAR-100 Classification

In this experiment, we leverage extra unlabeled data to further improve the CIFAR-100 classification accuracy. 3 different unlabeled tiny image datasets are adopted: (1) 50K training images of CIFAR-10; (2) 0.5M images which are selected from the 80M Tiny Image dataset [29]; (3) 1.3M images in Imagenet 2012 [23] training data after being down-sampled. The architecture we use is the network described in Section 4.1.

The results are shown in Table 3. When extra data is introduced, the accuracy shows a progressive increase as the size of the training data grows. It should be noticed that we achieve sate-of-the-art in a different way with other methods: we aim to find good initialization and use extra unlabeled data, while they tend to investigate delicate network architectures. Hence there is a potential to combine

		CIFAR100	CIFAR10
ELU [4]		75.72%	93.45%
Scalable Bayesian [25]		72.6%	93.63%
LSUV [19]		72.34%	94.16%
Resnet-v1-164/110 [10]		74.84%	93.39%
Resnet-v2-164 [11]		75.67%	93.57%
Resnet-v2-1001 [11]		**77.29%**	**95.08%**
Ours Pre-trained on	Data driven init	75.41%	94.93%
	CIFAR10 50K	75.68%	-
	Tiny 0.5M	75.82%	-
	Tiny Imagenet 1.3M	76.35%	95.08%

Table 3: Comparisons with state-of-the-arts: CIFAR100 and CIFAR10 classification accuracy (with simple data-augmentation). The architecture we use is a concise 15-layer fully convolutional network. Pre-trainings do not use any label.

our method with theirs to achieve better overall classification performance.

4.2.2 PASCAL VOC Classification

To better compare with prior work, we use the training images of Imagenet [23] 2012 without labels to learn feature by pre-training, and then evaluate the learned feature on PASCAL VOC. Yet we choose deeper VGG-16 based Network-2 in Table 1, instead of Alexnet [16] which is used by prior works. The reasons are: (1) Alexnet contains several severe down-sampling layers at first, which is quite harmful to our reconstruction pre-training, (2) the layers in Alexnet are carefully hand-crafted to ease the optimization, while our shortcut connections break the original data flow and make it hard to converge, and (3) for the encoder-decoder framework doubles the network depth, we want to verify that our method is suitable for deeper architectures such as 16-layer VGGNet.

We report our result as well as other Alexnet-based results in Table 4. Doersch *et al*. [5] and Pathak *et al*. [21] use unlabeled Imagenet data like ours, while Wang *et al*. [32] use extra unlabeled videos. Our improvement based on random initialized network is competitive.

Architecture	Pretraining	Accuracy
Alexnet-based	Random Gaussian	53.3%
	Doersch *et al*. [5]	55.3%
	Wang *et al*. [32]	**58.4%**
	Pathak *et al*. [21]	56.5%
Ours VGG-based	Random Gaussian	67.1%
	Ours	**71.3%**

Table 4: Comparisons with other CNN-based unsupervised methods on PASCAL VOC classification task.

4.2.3 STL-10 Classification

STL-10 is a dataset for unsupervised and semi-supervised learning. It contains 10 folds of labeled images with 1 thousand images each fold and 100 thousand extra unlabeled images. Network-1 in Table 1 with the parameters $m_1 = 3, m_2 = 3, m_3 = 4, m_4 = 5$ and $n = 64$ is used, which leads to a 30-layer auto-encoder for unsupervised pre-training and a 16-layer network for classification. The comparisons with other state-of-the-art methods are shown in Table 5. We achieve significant improvements against all the other methods. It should be noticed that the same classification network without using the unlabeled data leads to a very low accuracy of 59.81%, which indicates that our performance gain mainly comes from the use of the feature learned from unlabeled data.

Method	Accuracy
Swersky *et al*. [27]	70.1%
Target Coding [33]	73.15%
Huang *et al*. [12]	76.8%
Exemplar-CNN [6]	72.80%
Ours	**80.60%**

Table 5: Comparisons with state-of-the-arts: STL-10 classification accuracy.

4.2.4 Image Segmentation

In image classification, the decoder is discarded during fine-tuning. In this section we show that both the encoder and decoder are transferable. We evaluate our network on PASCAL VOC 2012 segmentation task. The existing CNN-based segmentation methods usually have encoder-decoder network architectures [20, 17, 1], which means that our auto-encoder network can be easily transformed for segmentation by replacing the final deconvolutional layer to convolution with suitable number of filters. Other CNN-based segmentation methods often use the pre-trained encoder part as initialization and initialize the decoder randomly. We train the segmentation network with 3 different initialization strategies: (1)random Gaussian numbers for all layers, (2)initialize the encoder by pre-training and initialize the decoder randomly, and (3) initialize both the encoder and decoder by pre-training. Pre-training is conducted on ImageNet 2012 training images without labels, and fine-tuning for segmentation is carried out on training images of PASCAL VOC 2012 segmentation task. The mean intersection over union scores on validation images are reported in Figure 4. It is easy to observe that the more layers we transferred from unsupervised pre-training, the better segmentation results we achieve.

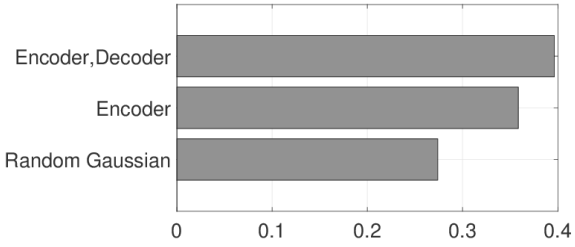

Figure 4: Mean IUs of PASCAL VOC 2012 Segmentation task on validation data with different fine-tuning strategies.

4.3. Discussions

Necessity of Shortcut Connections In this section we show that using multi-layer shortcut connections is crucial for reconstruction and feature learning. We train three Network-2 based auto-encoder on Imagenet 2012 training images with different shortcut strategies: (1) no shortcut connections, (2) only connect input/output and (3) our multi-layer shortcut connections. The reconstruction PSNR after pre-training and mean APs after fine-tuning on PASCAL VOC 2007 classification task are shown in Table 6. When using our shortcut strategy, both reconstruction performance and classification performance after fine-tuning are better.

	PSNR	mAP
No shortcut	22.19	68.94%
Input output shortcut	23.50	69.38%
Ours full shortcut	**23.87**	**71.25%**

Table 6: Comparions of different shortcut strategies on reconstruction PSNR and classification performance after fine-tuning.

Comparisons of Corruptions In section 3.3, we argued that different corruption types may affect the learned feature, which seems true from the observations of previous experiments: Pixel Gaussian noise is good enough for small images like CIFAR-10 and CIFAR-100, but small missing blocks performs better on larger images such as STL-10 and VOC. In this experiment, we compare different types of corruption mentioned in Section 3.3 on PASCAL VOC 2007 for both data-dependent initialization and unsupervised feature learning. One can see from Figure 2 for a illustration of the corruption. All other experimental settings are the same as those in Section 4.2.2.

The classification results are reported in Table 7. We can observe that pixel Gaussian noise is not a good choice in both cases. Patch Gaussian noise is good for data-dependent weight initialization, but performs worse than random initialization in unsupervised training. The reason may be that the reconstruction for this noise is too easy as the training data grows to 1.3M. The corruption of multiple small missing blocks outperforms others.

Corruption	Data driven initialization	Unsupervised pre-training
No pre-training	67.11%	
Pixel Gaussian noise	68.20%	69.23%
Patch Gaussian noise	69.44%	68.24%
Big missing block	68.25%	69.08%
Small missing blocks	**69.87%**	**71.25%**

Table 7: Compare different types of corruption: mAP for PASCAL VOC 2007 classification.

4.4. Understanding the Leaned Features

Although the results in previous sections show that our network can learn good features, it is still not clear how and what the network learns from unlabeled data. In this section, we firstly try to train our convolutional auto-encoder on unlabeled data with different distributions, and compare their influences on supervised fine-tuning. Then we provide some interesting visualizations to help understand the learned features.

4.4.1 Training on Images with Different Distributions

The 20 classes in PASCAL VOC dataset can be categorized into 4 coarse classes: animal, vehicle, indoor and person. In this section, we investigate whether pre-training on images with different distributions affects the overall recognition performance. For instance, if the network sees a lot of animal images during pre-training, we would like to know if it could better recognize animals.

Specifically, we take two subsets from Imagenet 2012 training images as our unlabeled datasets. Images of the first subset belong to super class 'conveyance, transport', and those of the second belong to 'mammal' or 'bird'. Each of them contain 93 thousand images. The two coarse classes instead of fine-grained classes are used because we assume that the training data is unlabeled. It is usually easier to get unlabeled data which belongs to a coarse category like 'animal' or 'vehicle' in practical.

The mean average precision score for each coarse class is shown in Table 8. We can observe that the vehicle 93K pre-trained network works better on recognizing vehicles and so does animal 93K on recognizing animals. Meanwhile, in each class, the classification networks with unsupervised pre-training outperform the one trained from scratch.

4.4.2 Feature Visualization

In this section, we try to analyse the features learned by our unsupervised learning method by visualization. We visual-

Pre-trained on	Animal	Vechicle	Mean All
No pre-training	67.02%	76.87%	67.11%
Animal 93K	**73.04%**	79.04%	70.77%
Vehicle 93K	72.17%	**80.00%**	**70.93%**

Table 8: Compare with different pre-training data: Mean average precision of PASCAL VOC 2007 classification on different coarse classes.

ize the 12th activation outputs of the pre-training network in Section 4.4.1, as shown in Figure 5. All images for visualization are not in training data.

(a) Original Images (b) Animal Trained (c) Vehicle Trained

Figure 5: Visualization of feature maps obtained by the denoising auto-encoders trained on different images. The network trained on animal images is more sensitive to dog faces, while the one trained on vehicle images is more sensitive to wheels. The network learns these merely by reconstructing corrupted images with different distributions, and no annotated label is used.

The animal trained network shows more attention on dog faces, while the vehicle trained network is more sensitive to wheels. It is intriguing because the networks learn these only by pixel-wise reconstruction, and there is no annotated label indicating where the dog or the car is in each training image. The reasons may be: (1) 'dog' and 'car' are dominant classes in each training data, and the network "remember" some typical patterns during reconstruction, and (2) complex image contents like dog faces and wheels are usually harder to reconstruct than other parts, hence the network gives them more attention.

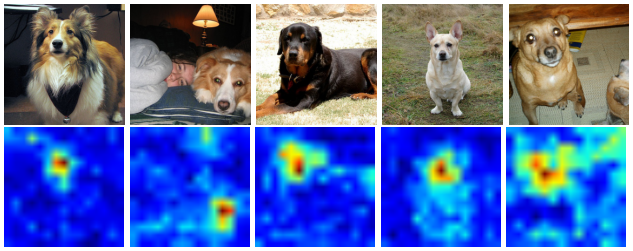

(a) Clean Images and Features

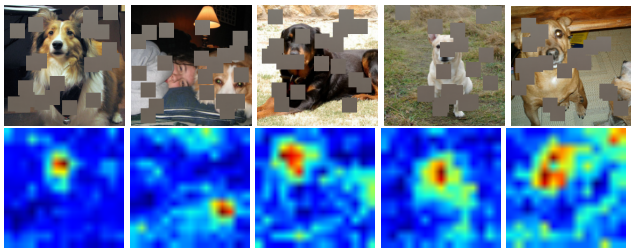

(b) Corrupted Images and Features

Figure 6: Visualization of more features of animal trained network. The features change little even when then dog faces are partially occluded.

We show more visualization of the animal trained auto-encoder in Figure 6a. It can act as a simple unsupervised trained dog face detector in these cases, although the training images is not all dogs but with other kinds of animals. We further show the features of corrupted images in Figure 6b. One can see that the learned features change little even when the dog faces are partially occluded, which shows that the features learned by our denoising auto-encoder are very robust to corruption.

5. Conclusions

A convolutional denoising auto-encoder for unsupervised feature learning and data driven network initialization is proposed in this paper. We hit the state-of-the-arts CIFAR-10 and CIFAR-100 classification merely through our initialization strategy with a quite concise network architecture. The classification performance on CIFAR-100 are further improved by making use of more extra unlabeled data. When the amount of labeled data is limited, our network show even more potential. We also set new records on STL-10 dataset by a huge margin compared to other methods. Experimental results show that the learned feature can be easily transferred to different high level vision tasks such as image classification and semantic segmentation.

References

[1] V. Badrinarayanan, A. Kendall, and R. Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *arXiv preprint arXiv:1511.00561*, 2015.

[2] Y. Bengio. Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, 2(1):1–127, 2009.

[3] T. Chen, M. Li, Y. Li, M. Lin, N. Wang, M. Wang, T. Xiao, B. Xu, C. Zhang, and Z. Zhang. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *NIPS Workshop on Machine Learning Systems*, 2016.

[4] D. Clevert, T. Unterthiner, and S. Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). In *Proc. Int. Conf. Learn. Representations*, 2016.

[5] C. Doersch, A. Gupta, and A. A. Efros. Unsupervised visual representation learning by context prediction. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 1422–1430, 2015.

[6] A. Dosovitskiy, J. T. Springenberg, M. A. Riedmiller, and T. Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Proc. Advances in Neural Inf. Process. Syst.*, pages 766–774, 2014.

[7] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proc. Int. Conf. Artificial Intell. & Stat.*, pages 249–256, 2010.

[8] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio. Generative adversarial nets. In *Proc. Advances in Neural Inf. Process. Syst.*, pages 2672–2680, 2014.

[9] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 1026–1034, 2015.

[10] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016.

[11] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In *Proc. Eur. Conf. Comp. Vis.*, pages 630–645, 2016.

[12] C. Huang, C. Change Loy, and X. Tang. Unsupervised learning of discriminative attributes and visual representations. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 5175–5184, 2016.

[13] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proc. Int. Conf. Mach. Learn.*, pages 448–456, 2015.

[14] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *Proc. Int. Conf. Learn. Representations*, 2015.

[15] P. Krähenbühl, C. Doersch, J. Donahue, and T. Darrell. Data-dependent initializations of convolutional neural networks. In *Proc. Int. Conf. Learn. Representations*, 2016.

[16] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proc. Advances in Neural Inf. Process. Syst.*, pages 1106–1114, 2012.

[17] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 3431–3440, 2015.

[18] X. Mao, C. Shen, and Y. Yang. Image restoration using convolutional auto-encoders with symmetric skip connections. In *Proc. Advances in Neural Inf. Process. Syst.*, 2016.

[19] D. Mishkin and J. Matas. All you need is a good init. In *Proc. Int. Conf. Learn. Representations*, 2016.

[20] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 1520–1528, 2015.

[21] D. Pathak, P. Krähenbühl, J. Donahue, T. Darrell, and A. A. Efros. Context encoders: Feature learning by inpainting. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2016.

[22] M. Ranzato, F. J. Huang, Y. Boureau, and Y. LeCun. Unsupervised learning of invariant feature hierarchies with applications to object recognition. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2007.

[23] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and F. Li. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vision*, 115(3):211–252, 2015.

[24] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proc. Int. Conf. Learn. Representations*, 2015.

[25] J. Snoek, O. Rippel, K. Swersky, R. Kiros, N. Satish, N. Sundaram, M. M. A. Patwary, Prabhat, and R. P. Adams. Scalable bayesian optimization using deep neural networks. In *Proc. Int. Conf. Mach. Learn.*, pages 2171–2180, 2015.

[26] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. A. Riedmiller. Striving for simplicity: The all convolutional net. In *Proc. Int. Conf. Learn. Representations*, 2015.

[27] K. Swersky, J. Snoek, and R. P. Adams. Multi-task bayesian optimization. In *Proc. Advances in Neural Inf. Process. Syst.*, pages 2004–2012, 2013.

[28] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 1–9, 2015.

[29] A. Torralba, R. Fergus, and W. T. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(11):1958–1970, 2008.

[30] H. Valpola. From neural PCA to deep unsupervised learning. *Advances in Independent Component Analysis and Learning Machines*, pages 143–171, 2015.

[31] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P. Manzagol. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *J. Mach. Learn. Res.*, 11:3371–3408, 2010.

[32] X. Wang and A. Gupta. Unsupervised learning of visual representations using videos. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 2794–2802, 2015.

[33] S. Yang, P. Luo, C. C. Loy, K. W. Shum, and X. Tang. Deep representation learning with target coding. In *Proc. AAAI Conf. Artificial Intell.*, pages 3848–3854, 2015.

Supplemental Material

A. Corruption Parameters

In this section we describe the corruption parameters we use in detail. In general, we find that the influence of the corruption parameters on the final results is small. For pixel Gaussian noise and patch Gaussian noise, we choose $\sigma \in [10, 50]$ uniformly and randomly for each image. Although traditional denoising problem normally assumes a fixed σ, the CNN based denoising network can perform blind denoising. More importantly, we observe that using a random σ outperforms using a fixed one in most cases. For block mask corruption, assuming the side length of the fed image is l, the side length is taken randomly in $[l/4, l/2]$ for big block and $l/8$ for small blocks. The number of blocks for small block mask is determined by the number of unlabeled training images n. We set it as 21 for $n < 100,000$ and 35 otherwise.

B. The Number of Labeled Images

In Section 4.2.3, we report the results using pre-defined 10-fold training on STL-10. We explore the influence of the number of labeled images during fine-tuning on SLT-10 in this section. The results are shown in Figure 7. The main observation is that when using unlabeled images for pre-training, the network needs far less labeled ones to reach the same accuracy. Meanwhile, the less labeled images are used, the more improvement our method achieves comparing to the network without using unsupervised pre-training.

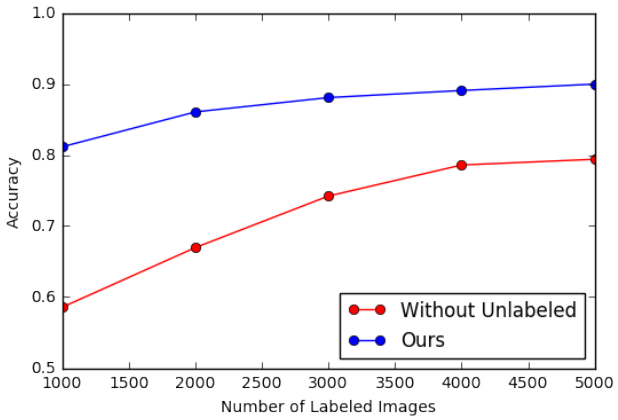

Figure 7: STL-10 testing accuracy on different number of labeled images.

C. Reconstruction Performance

Although our ultimate purpose is to learn better features for high level vision tasks instead of simply improving the reconstruction performance, we do observe that as other parameters being fixed, a better reconstruction performance always leads to a higher accuracy after fine-tuning. We show more reconstruction results on different datasets in Figure 8 9 10.

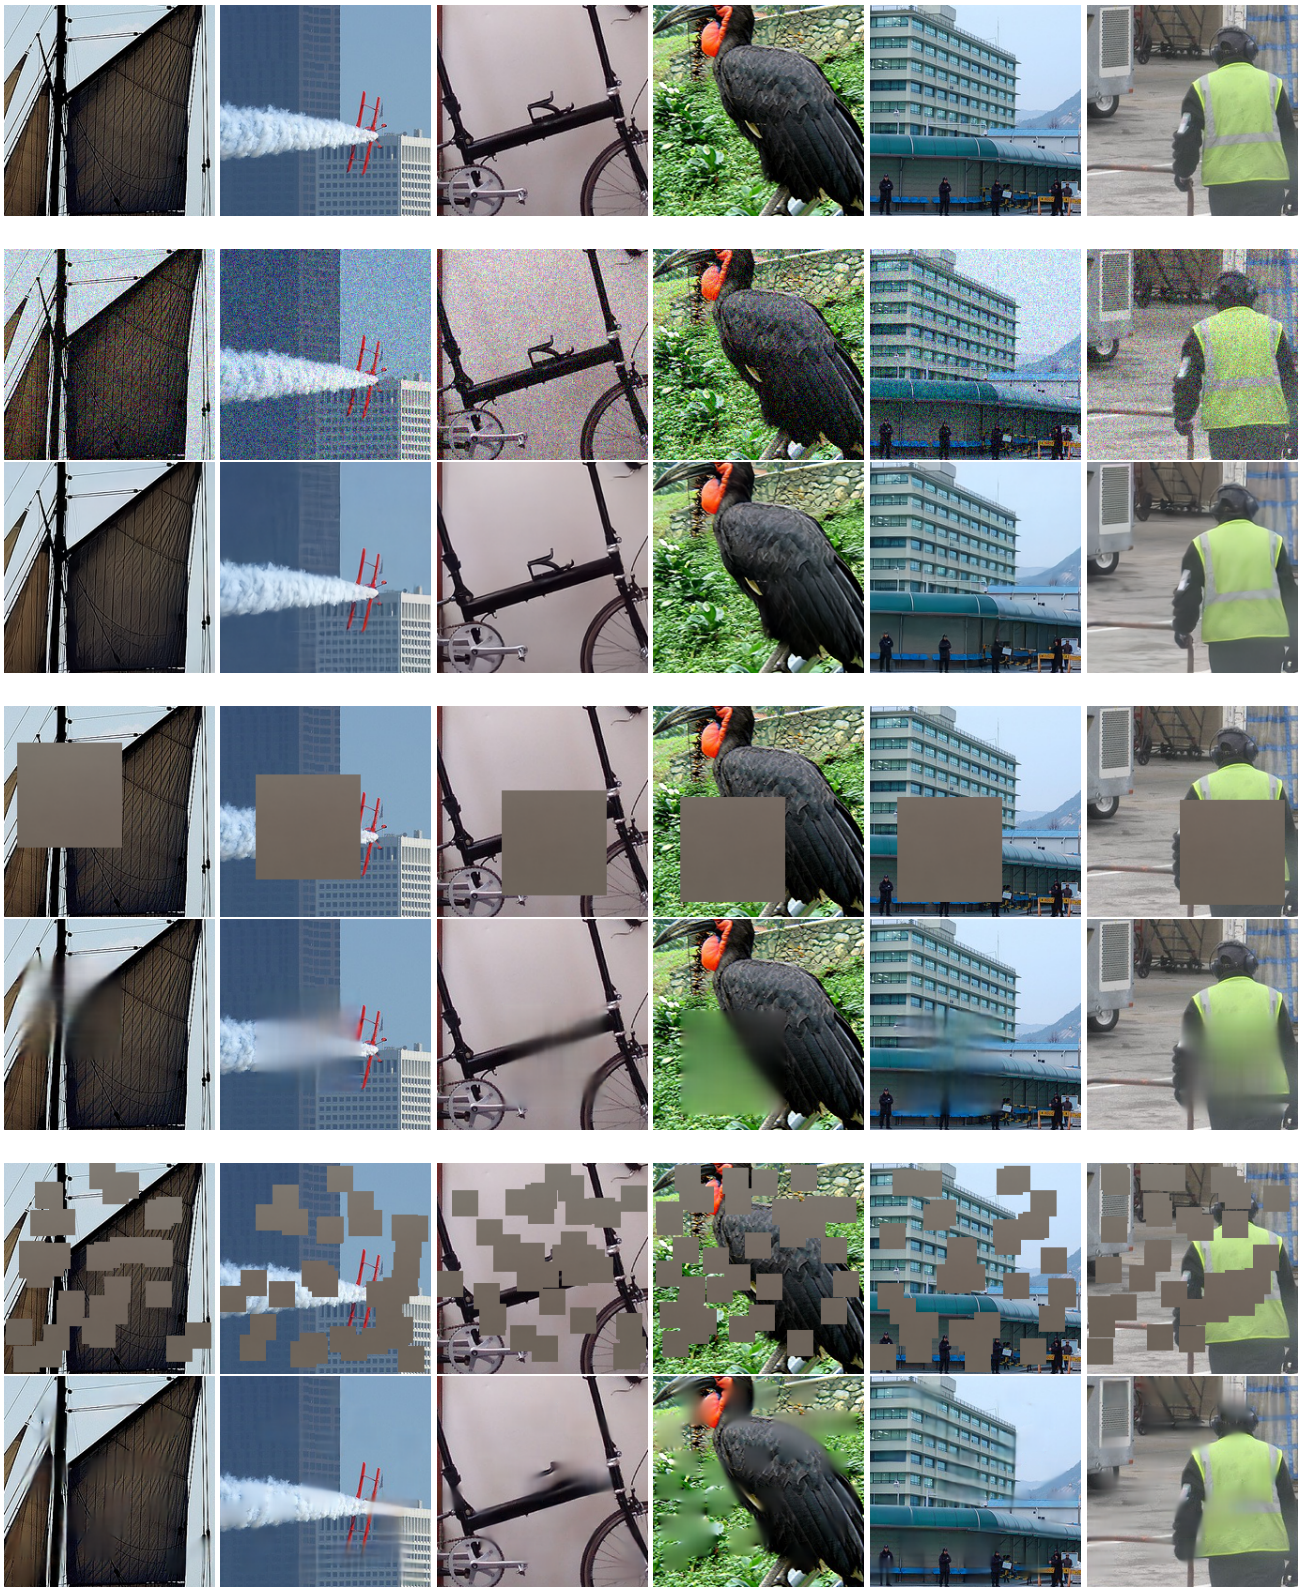

Figure 8: Reconstruction on PASCAL VOC images. The training is conducted on training data of Imagenet 2012. Images at the top line are clean ones.

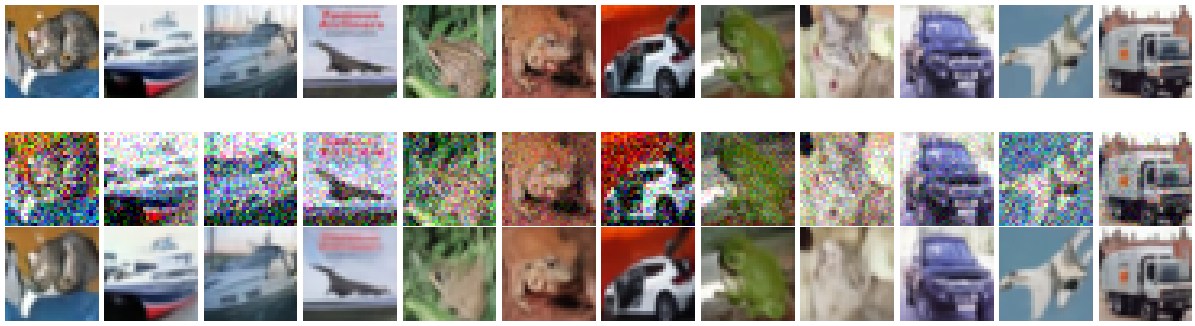

Figure 9: Reconstruction on CIFAR-10 testing images. The training is conducted on training data of CIFAR-10. Images at the top line are clean ones.

Figure 10: Reconstruction on STL-10 testing images. The training is conducted on 100 thousand unlabeled images of STL-10. Images at the top line are clean ones.